

PL-TR-96-2130

A REGIONAL WEATHER FORECASTING EXPERT SYSTEM

**William H. Jaspersen
David E. Venne**

**Augsburg College
Center for Atmospheric & Space Sciences
2211 Riverside Avenue
Minneapolis, MN 55454**

15 June 1996

19961018 118

**Final Report
15 December 1991-15 January 1996**


APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED





**PHILLIPS LABORATORY
Directorate of Geophysics
AIR FORCE MATERIEL COMMAND
HANSCOM AFB, MA 01731-3010**

DTIC QUALITY INSPECTED 3

"This technical report has been reviewed and is approved for publication."


ROSEMARY M. DYER
Contract Manager


DONALD C. NORQUIST
Acting Chief, Satellite Analysis and Weather
Prediction Branch
Atmospheric Sciences Division


DONALD A. CHISHOLM, Acting Director
Atmospheric Sciences Division

This report has been reviewed by the ESC Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS).

Qualified requestors may obtain additional copies from the Defense Technical Information Center (DTIC). All others should apply to the National Technical Information Service (NTIS).

If your address has changed, or if you wish to be removed from the mailing list, or if the addressee is no longer employed by your organization, please notify PL/IM, 29 Randolph Road, Hanscom AFB, MA 01731-3010. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 15 June 1996	3. REPORT TYPE AND DATES COVERED Final (15 December 1991 - 15 January 1996)		
4. TITLE AND SUBTITLE A Regional Weather Forecasting Expert System		5. FUNDING NUMBERS PE 63707F PR 2688 TA 01 WU LA Contract F19628-92-K-0002		
6. AUTHOR(S) William H. Jasperson David E. Venne				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Augsburg College Center for Atmospheric and Space Sciences 2211 Riverside Ave. Minneapolis, MN 55454		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Phillips Laboratory 29 Randolph Road Hanscom AFB, MA 01731-3010 Contract Manager: Rosemary Dyer/GPAB		10. SPONSORING/MONITORING AGENCY REPORT NUMBER PL-TR-96-2130		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report describes a prototype weather forecasting expert system called <i>Itasca</i> that provides a short-term (0 - 12 hour) weather forecast for a midlatitude station when only limited data are available. If data are available only at the forecast station, single-station forecasting techniques are utilized. If more stations are available, their data are used to improve the forecast. The expert system simulates the process used by experienced human forecasters by first constructing a meteorological analysis and then, from this analysis, making a forecast. The system's user interface consists of pull-down menus in a windowed environment. The system presents data graphically as much as possible.				
14. SUBJECT TERMS Expert System, Meteorology, Forecasting		15. NUMBER OF PAGES 22		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

TABLE OF CONTENTS

1.	INTRODUCTION	1
2.	REPRESENTATIONAL STRUCTURES	2
2.1	Definitions	2
2.2	Example from <i>Itasca</i>	3
2.2.1	Stations and Observations	3
2.2.2	Meteorological Entities	5
2.2.2.1	High and Low Pressure Centers	5
2.2.2.2	Fronts	8
3.	ANALYSIS	10
3.1	Gridded Fields	10
3.2	New Analysis	11
3.2.1	Low Pressure Centers	11
3.2.2	Fronts	11
3.2.3	High Pressure Centers	12
3.3	Update Analysis	12
4.	FORECASTS	14
5.	USER INTERFACE	16
6.	SUMMARY	18

1. INTRODUCTION

The goal of this project is to develop a meteorological forecasting system that can be used in the field for the purpose of providing short-term weather forecasts for regions in the middle-latitudes. The forecasting system is a knowledge-based system in that it incorporates knowledge about meteorology and the relationships between meteorological variables, so called textbook knowledge, as well as "rule-of-thumb" knowledge acquired by experienced weather forecasters.

For the military, "first-in" capability is of vital importance, but during this time the availability of data is often very limited. A useful system for this environment must not be dependent upon extensive data or even a prescribed set of data. As time moves forward, data may become available at more stations and data may become unavailable at stations that had previously provided data. The system must use the available data in an opportunistic manner such that it makes the best short-term forecast that it can with the information that it has.

A second desirable feature of the system is that it should, in addition to a forecast, provide information to the user such that if the user has additional knowledge, he or she will be able to confirm or to add quality to the forecast. In other words, the system can be used in a forecaster assisting mode. At an intermediate level this assistance can be provided in the form of data displays, maps and diagrams that provide data visually to the forecaster in a form to which they are accustomed. At an advanced stage, the user may have knowledge that the system does not have or have a different interpretation of the data that it does have, and they should be able to interact to change the system's analysis in such a way that the forecast can be better.

In order that this type of interaction with the user is possible, the system under development is separated into two separate and distinct functions. First, the system takes all data available to it and constructs a synoptic environment. This is the analysis phase of the forecasting process. Second, upon the user's command, the system will produce a forecast based upon this analysis of the synoptic environment and the observations available to it.

This report describes the representational structures and the process that this knowledge-based system, called *Itasca* in this report, uses to produce short-term forecasts. The following section defines the required terminology and describes the primary representations used in *Itasca*. Sections 3 and 4 describe the analysis and forecasting process, respectively, and Section 5 describes the user interface. Section 6 provides a summary.

2. REPRESENTATIONAL STRUCTURES

The software development environment used in the construction of *Itasca* is Neuron Data's Smart Elements product. Smart Elements consists of an interface development tool called Open Interface and an expert system development tool called Nexpert-Object. Open Interface allows the construction of user interfaces that are transportable over a wide range of hardware platforms. Nexpert-Object supports a broad range of representation features and capabilities in the development of knowledge based systems. Before describing the analysis and forecasting process used in *Itasca*, it is necessary to define and review the basic representational structures that are used to represent the observational data and the meteorological entities created by the system.

2.1 Definitions

An *object* is the smallest "chunk" of information that is contained in the knowledge-based system. While an object may represent a single variable, it is usually used to represent a more complex entity. In *Itasca*, for example, meteorological features such as high and low pressure systems and fronts, meteorological data such as surface and upper-air observations, geographic entities such as meteorological observation stations are each represented as an object.

A synoptic map may consist of multiple fronts, multiple stations, etc. It makes sense that like entities, all of whom share common features, should be grouped. This grouping is accomplished by defining a *class* which contains all of the objects that are a part of that class. Then all of the fronts that appear on a map are objects of the class "fronts." All of the high pressure systems are members of the class "high_pressure_centers," and so forth.

Classes may also be *subclasses* of a higher level class, or *parent class*. For example, in *Itasca* the classes "fronts," and "pressure_center," are both subclasses of the parent class "meteorological_entities." And the classes "high_pressure_center" and "low_pressure_center" are subclasses of the parent class "pressure_center." As can be seen from the above example, a subclass is a class which represents a subset or specialization of another class. Using the example above, a "front" is quite different from a "pressure_center," but both are categorized as "meteorological_entities." In addition, any class may be a subclass of multiple parent classes.

Just as some classes may be defined as subclasses of parent classes, N-O supports the representation of objects that may be defined as *subobjects* of other objects. The relationship of a subobject to an object can be thought of as a relationship of the type "is a part of" or "is associated with."

In general, each class or object has *properties* (sometimes called *slots*) associated with it that are used to define that particular class or object. The *values* of the properties distinguish objects in a class from one another.

One of the important reasons for developing a complete class and object structure is that these properties and the methods for finding values for these properties may be inherited down the class/object structure. This inheritance capability allows the properties to be defined at the class level and to flow down to the appropriate object level. When an object, a low pressure center for example, is created, it inherits all of the properties that have been defined in its parent class. All low pressure centers, for instance, have properties of speed, direction of motion, and central pressure as well as many other properties. One of the functions of the knowledge base is to provide, for each of these objects, values for its properties.

Another important reason for a complete class and object structure is that Nexpert-Object provides mechanisms to reference and process all objects that are members of a given class or all objects that are subobjects of another object. This will be illustrated in the following examples.

2.2 Examples from *Itasca*

This section presents a description of the observation and meteorological classes and objects used in *Itasca*.

2.2.1 Stations and Observations

A station object is created when a station first reports a meteorological observation. This station object is created as a member of the class "station" which is a subclass of a geometrical class called "point." The class "point" has the geometric properties of latitude, longitude and elevation. The class "station" inherits the geometric properties of its parent class "point" and adds identification properties such as its WMO and ICAO number/name, an arbitrary station identification number and name (used internally), and the number of hours the station is offset from Greenwich Mean Time. When an object is created for a new station, it inherits the properties of all of its parent classes, and the appropriate values are assigned to each of its properties. The properties of class "station" are time independent, so the values of each object remain the same over time.

Meteorological surface and upper-air observations are also defined through a class/object structure. The class "sfc_ob" has properties for the observed meteorological variables of temperature, dew point temperature, wind direction and speed, pressure, visibility, weather, cloud type, height, and amount for up to four levels, total sky cover, and opaque sky cover. Some derived variables such as pressure tendency and low, middle, and high cloud amounts, and the local time are also defined "sfc_ob" properties. The class "upa_ob" has properties for variables that are computed from the sounding such as stability indices, precipitable water, low level advection, lifting condensation level, and the convective temperature. Values of temperature, wind speed and direction, and humidity are not maintained within the class/object structure but may be obtained at pressure levels through function calls to the *Itasca* database. Meteorological observations

also include the identification properties of its object name, its object type (e.g. *sfc_ob* or *upa_ob*), and the name of its station object and the unique identification number representing the station at which the observation was taken.

The observation classes in *Itasca* are all subclasses of the class "time_dependent_object." This class contains properties of time such as year, month, day, and hour, as well as the property "hrs_from_present" which represents the number of hours from the present that the observation was taken. All of these properties are inherited by time dependent objects at the time they are created.

Each hourly cycle of *Itasca* results in new data becoming available for one or more surface/upper-air stations. A new surface and/or upper-air object is created for each of these new observations. If data are received from a station that had not previously reported, a station object is first created for that station. The value of the property "hrs_from_present" for meteorological observations (and all objects that descend from the "time_dependent_object" class) are incremented by one hour, and the new observation objects are given the value of zero for this property. Observation objects are deleted when they become 48 hours old, but the data remain in the *Itasca* database.

Figure 1 illustrates the relationships between the station and observation classes and the objects created from them. The circles indicate the Station, Surface Observation, and the Upper-Air Observation classes, of which the latter two are subclasses of the Time Dependent Object class. In this example, two stations exist and each has provided two surface observations. Station A has also provided an upper-air observation. Each station is an object of the class "station," and each observation is an object of the appropriate observation class as well as a subobject of the appropriate station object.

If ten stations report hourly, 480 surface observation objects will exist after two days. This structure of classes and subclasses allows *Itasca* to easily obtain the data it requires for a particular process. For example, all of the current surface observations are obtained by requesting all members of the Surface Observation class with a value for the property "hrs_from_present" equal to zero. Or, all of the surface observations for station A may be found by taking all of the subobjects of station A that have the type "sfc_ob."

This has been a simplified discussion of the station/observation class and object structure. In *Itasca*, there are two other types of objects besides observation objects that are attached to station objects. The first set of objects contain topographic information about the station. Information includes the direction, distance and extent of water bodies and terrain features (flat, hilly, rugged) from the station. These objects are created once and do not change. The second set of objects are station environment objects and are created during the analysis process. They contain the names of the closest fronts and high and low pressure systems to the station, and the time of the most recent frontal passage. These objects are time dependent objects and new objects are created each hour and retained for six hours.

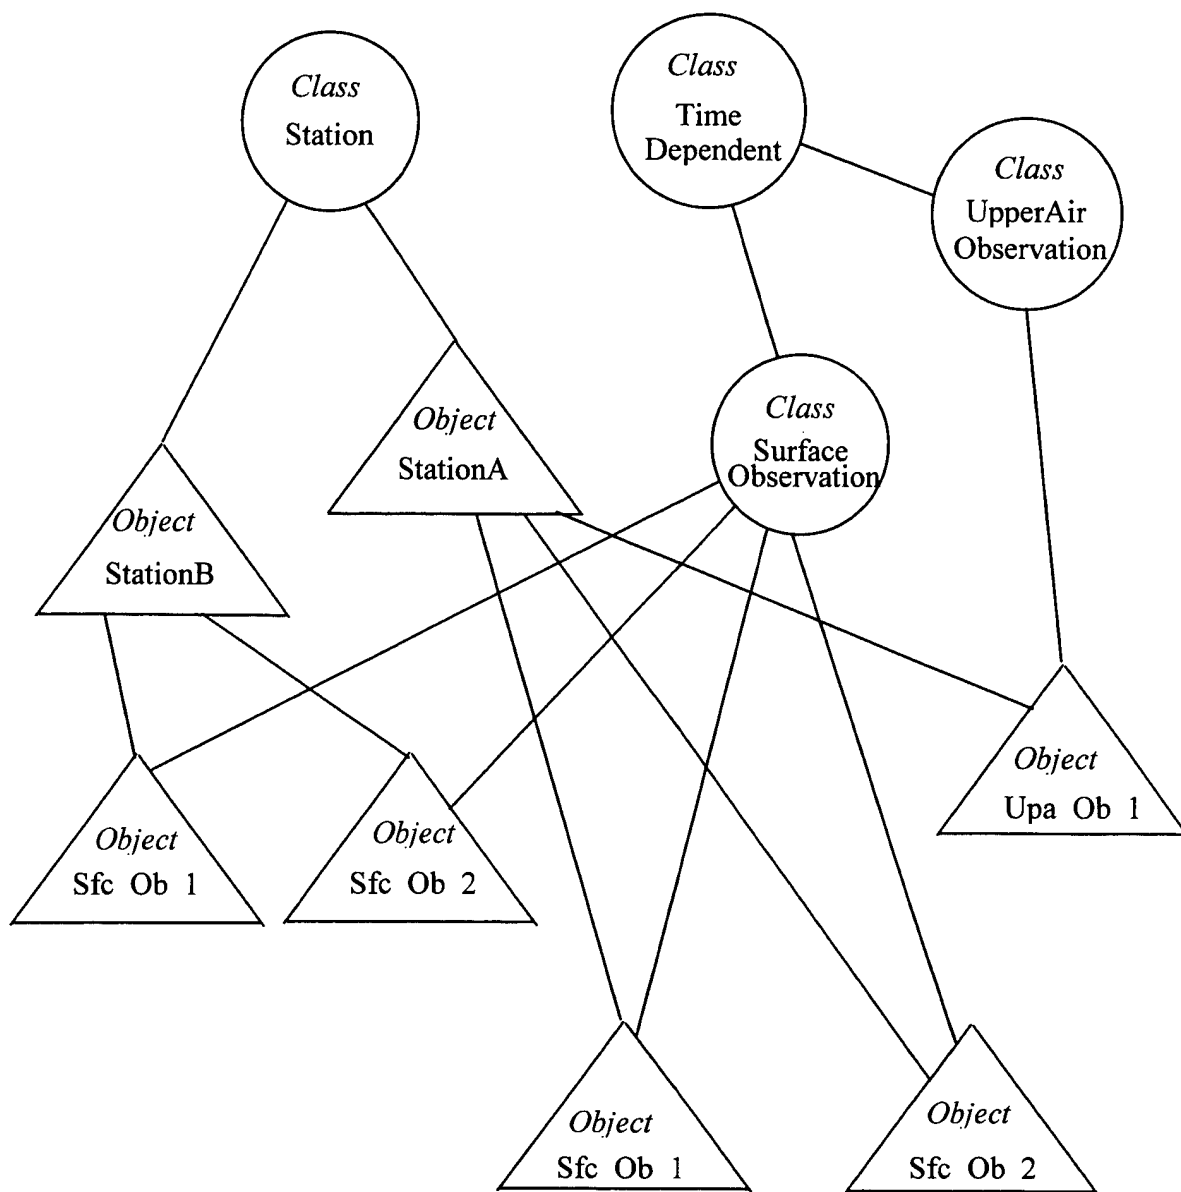


Figure 1: Station and Observation Classes and Objects.

2.2.2 Meteorological Entities

The purpose of the analysis phase of *Itasca* is to generate a synoptic map. This synoptic map consists of high and low pressure systems and fronts that are inferred from the observational data as well as their attributes such as position, speed, movement, etc. This section describes the class/object structure of these meteorological entities.

2.2.2.1 High and Low Pressure Centers

The classes "high_pressure_center" and "low_pressure_center" are both subclasses of the class "pressure_center" which is a subclass of "met_entity." Since all meteorological entities exhibit time dependent behavior, they are also subclasses of the class "time_dependent_objects."

When the analysis indicates the need for a new high pressure center, two objects are created. The first is a high pressure center supervisor. This object is a member of the class "high_super" which is a subclass of the class "met_entity_supervisor." There are as many high pressure center supervisor objects as there are highs in the synoptic analysis. The second object created represents the high pressure center itself. It is created as a member of the class "high_pressure_center" as well as a subobject of its high pressure center supervisor object.

The high pressure center supervisor has the properties of name (the name of the supervisor object), type (a "high pressure" supervisor), age (the number of hours since its creation), and up_to_date_name (the name of the high object at the current time).

The properties of the high pressure center object include its name, latitude and longitude of its center, central pressure, speed, direction of motion, and the name of its supervisor object. This object also has the property "hrs_from_present" whose value is set to zero when the high is first created. During each hourly cycle of *Itasca*, the value of the "hrs_from_present" property is incremented by one hour, a new high_pressure_center object is created and its name is put into the value of the up_to_date_name property of its supervisor and the values of its other properties are updated. Also the value of the age property of the supervisor is incremented by one.

Figure 2 illustrates the relationships between the classes and objects for high pressure centers. The circles indicate the Met_Entity, Pressure_Center, High_Pressure_Center, Met_Entity_Supervisor, and High_Super classes. In this example, there are two high pressure centers, A and B, each with its supervisor. Since there are two objects for both High_A and High_B, each has existed for two hours. All Met_Entity objects with hrs_from_present greater than 24 are deleted each hour. If the high is no longer present in the data, all of the high_pressure_center objects and their supervisor are deleted.

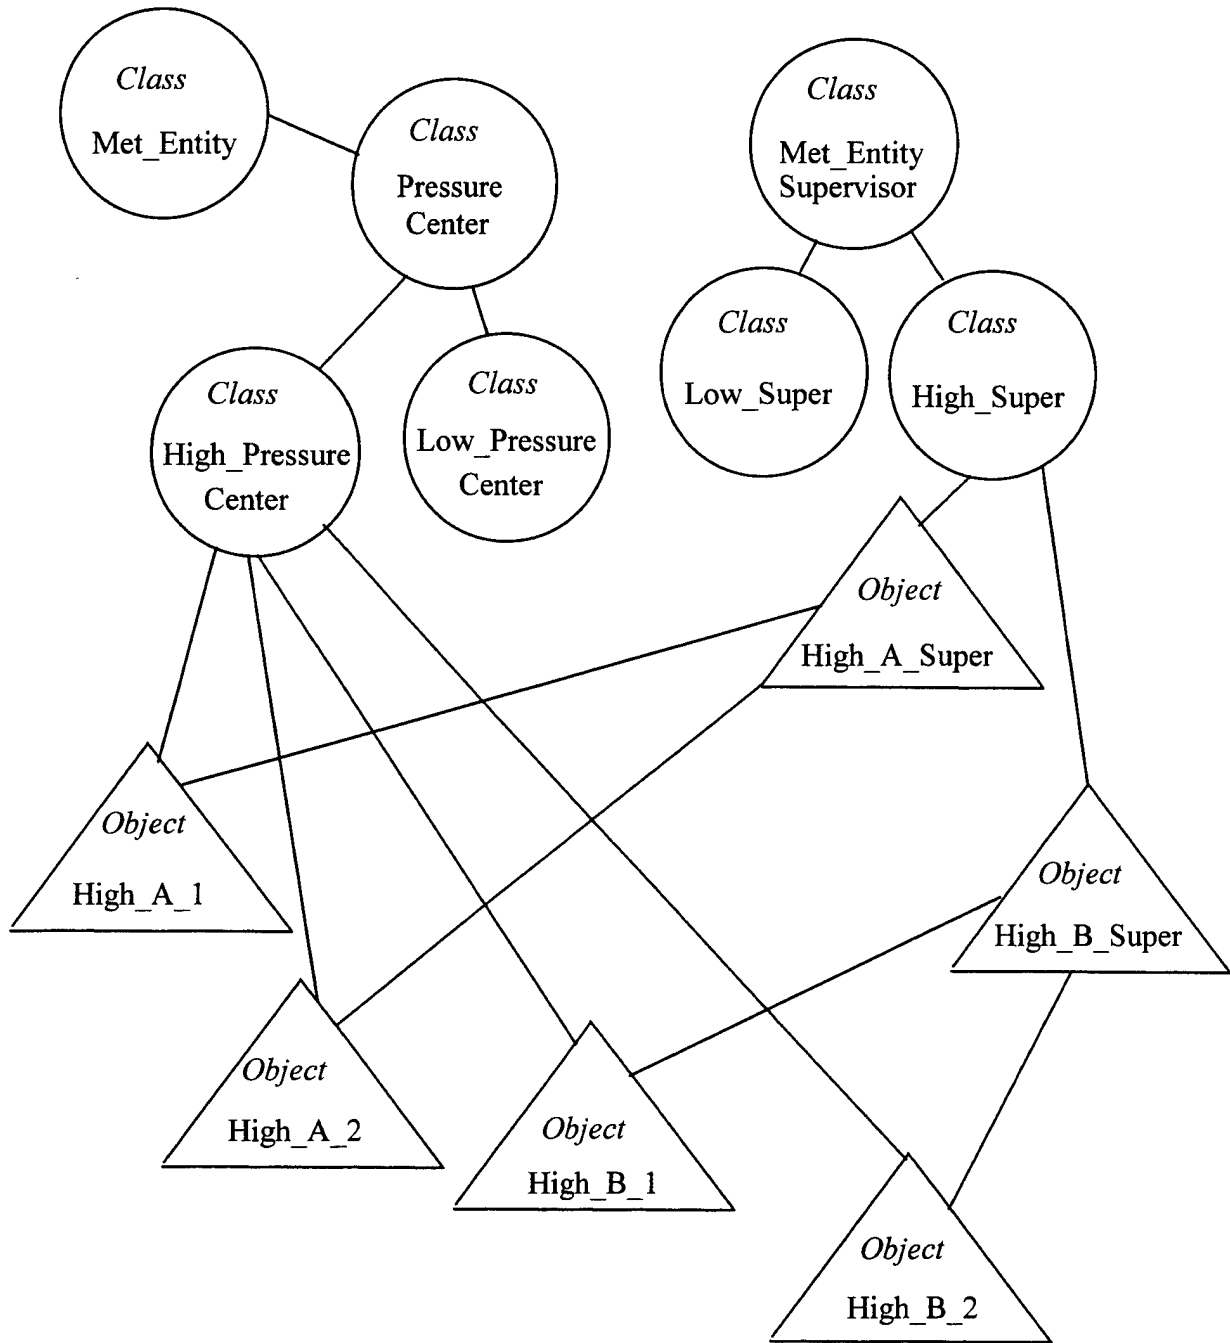


Figure 2: Class/Object Structure of High Pressure Centers

Low pressure centers have an entirely equivalent representation. The class "low_pressure_center" is a subclass of "pressure_center" and the class "low_super" is a subclass of "met_entity_supervisor." Each low pressure center is represented by a supervisor object that has properties of name, type, age and up_to_date_name, and each hour's low pressure center object belongs to the class "low_pressure_center" and is a subobject of its supervisor object.

All objects of current lows can be obtained by finding all of the objects of the class "low_pressure_center" for which the value of the property "hrs_from_present" is equal to zero. Hour by hour realizations of any given high or low can be obtained by collecting all of the subobjects of the supervisor. Maintaining this history of the highs and lows has advantages. For example, past analyses can essentially be recovered if short term fluctuations in the data result in an improper analysis for recent hours. Secondly, in some cases the position of highs or lows are known well enough over time that speed and direction of motion may be calculated directly from changes in position of the entity.

2.2.2.2 Fronts

The differences between fronts and the highs and lows already described arise from differences in the properties that are associated with fronts. High and low pressure centers are points, while fronts are lines. In *Itasca*, fronts are represented by a set of points called "front_points." Practically, the number of front_points used to describe a front ranges from two to ten. If there are data from only one station, the front will be represented by two or three points. If there are data from many stations, typically four to eight points will be used to represent the front. These points are determined by the knowledge base during the analysis phase. Each front point is an object with properties including a latitude, a longitude, a speed, and whether or not it is an endpoint. Once determined, the front_points are maintained within the *Itasca* database, not in the front object itself. *Itasca* constructs a line that passes through the front_points and can graphically present the front through the user interface. Also, questions in the analysis or forecasting knowledge base that require the distance to, the orientation of, or the speed of the front are provided by *Itasca* using the mathematically generated frontal curve and not, for example, the closest front_point.

Apart from the fact that the representation of front positions are maintained within a database rather than the front objects, fronts are handled in an analogous way to the high and low pressure centers. The class "front" is a subclass of "met_entity," and for each front there is a class "front_super" that is a subclass of "met_entity_supervisor." Each front object is a member of the class "front" as well as a subobject of its "front_super" supervisor object. Each hour each front object property "hrs_from_present" is incremented, another object is created, and the age and up_to_date_name of the front supervisor are updated.

There are a number of properties of the class "front" that are different from the high and low pressure centers. There is a property "type" (cold, warm, stationary, occluded) and a property of "orientation" (angle of the front). There are also properties that define the front's relationship to the other meteorological entities. Values for the properties end1 and end2 specify the name of any meteorological object to which the front is attached. For example, most fronts begin at a low pressure center. If this is so, the value of end1 (or end2) will be the name of that low pressure center object. If it is attached to another front, the value will be the name of that front object. Properties of end1_type and end2_type indicate the type of meteorological entity that is the endpoint (i.e. low pressure center or front). If no meteorological entity constitutes an endpoint, the value of the property is left blank. Three more properties called "airmass_left," "airmass_right," and "airmass_toward" are used to define the position of the front relative to the high pressure systems and direction of movement of the front. The airmasses are the high pressure systems that are on each side of the front. The property "airmass_left" will take the value of the name of the high pressure center of one side of the front and the property "airmass_right" will take the value of the name of the high pressure center on the other side of the front. The value of the property "airmass_toward" will be the name of the high pressure center that the front is moving toward.

Another difference with the front objects is that besides being objects of their front class and subobjects of their supervisor object, they are also created as subobjects of the low pressure center object to which they are attached. This construction allows multiple fronts to be attached to low pressure centers without having to keep track of any counters or prespecifying properties in low pressure center objects for names for an arbitrary number of fronts that are attached to it. All fronts attached to a low pressure center can be determined by making a list of all subobjects of the low pressure center object of type "front."

3. ANALYSIS

The purpose of making a meteorological analysis is to provide a physical interpretation for the weather occurring at a station and a physical basis upon which to base a short-term weather forecast. The analysis process consists of constructing a synoptic environment with the presence, location, and movement of pressure systems and fronts.

Itasca is designed to provide an analysis and forecast from an arbitrary number of one or more data reporting stations. The quality of the analysis depends largely on the number and distribution of observation stations. Because data do not have to be reported every hour, *Itasca* grids data whenever possible.

3.1 Gridded Fields

As part of the setup process for *Itasca*, a region that contains all of the stations that may provide data is defined along with an appropriate grid spacing. Whenever possible, *Itasca* generates gridded fields of variables. Besides the fact that stations may not report data on an hourly basis, advantages of gridding data include easy computation of gradients and display of fields in a manner consistent with a meteorologist's experience. Also this allows for the future implementation of providing initial data fields for system startup, if they are available.

Normally, data must be available from three or more non-colinear stations to produce a useful gridded field. An exception to this is a pressure field which may be constructed from one or two stations provided that the pressure gradient can be properly inferred from the wind direction. The quality of the wind direction is very important in constructing the analysis, particularly when there are only one or two stations. A wind quality knowledge base assigns a "quality factor" to the wind at the start of the analysis process. This quality factor is based upon the strength of the wind, time continuity (if any), spatial consistency (if any), current weather and local topography. If this quality factor is greater than a specified threshold, the wind direction is used to enhance the gridded pressure field, and therefore, the pressure analysis.

Within any gridded field, the reliability of the data is best within the area of observation. If there are three or more non-colinear stations, *Itasca* constructs a polygon that encloses all of the reporting stations and has the requirement that its perimeter is as small as possible. This polygon is referred to in this report as the "hull." Within the hull, gridded values are considered reliable; outside the hull, gridded values become less reliable as the distance to the hull increases. Note that the hull for two stations (or more colinear stations) is a line and for a single station is a point.

3.2 New Analysis

The process involved in making the analysis is somewhat different if the system is making a new analysis versus updating an analysis from the prior hour. This section provides a brief description of how a new analysis is made. In a new analysis, *Itasca* first creates low pressure centers, then fronts, and finally high pressure centers.

3.2.1 Low Pressure Centers

If a hull that encloses an area exists (i.e. there are three or more non-colinear stations), the analysis process begins checking if there is a local low pressure value within the hull. If so, a low pressure center is created at the point of low pressure. Values for speed and direction of motion are estimated from the 500 mb wind above the point if sounding data are available. Otherwise default values are used. If a low pressure center is not in the hull, as is typically the case, all points on the hull that have a local pressure minimum are found. Because *Itasca* assumes it has data from a relatively confined region (1000 km or so), there are rarely more than two pressure minima on the hull. Two or more minima usually corresponds to a front passing through the hull.

All local hull minima with an associated pressure gradient directed outward indicate a low pressure center outside the hull. The analysis knowledge base estimates the location of the low pressure center from the direction and shape of the pressure gradient at and near the hull local minimum and observation data from stations near the pressure minimum. If there is another local hull minimum with an outward directed gradient (not close to the first) another low will be created.

If there are only one or two stations, any low pressure center that exists must lie outside the hull. The location of any low pressure center can only be estimated by using observational data from the station or stations. If the wind direction does not have sufficient quality associated with it and there are no past data to provide a pressure tendency, no analysis will be made at this time and the resulting forecast will be based on climatologically modified persistence.

3.2.2 Fronts

A cold front is usually associated with a low pressure center. If there is only one station or if there is more than one station and there is no evidence of a front passing through the hull, a cold front is created with an orientation consistent with the pressure gradient and the wind field ahead of or behind the front.

Cold fronts that pass through the hull are usually reflected both in station data and the presence of a trough in the pressure field. Because the distribution of observational data may not be uniform, the best placement of the front is determined by first creating front points that lie near the pressure trough and then, if necessary, modifying their

location using observation data. Computationally, finding a trough with an arbitrary orientation is a difficult task and sometimes results in a misplaced front.

Warm fronts often do not have well defined troughs associated with them and have to be inferred from winds, temperature, temperature gradient, advection, and upper-air data.

Front_points that are defined within the hull have positions that may be relatively well known. The speed of a front as it passes through a hull can then be well determined as a function of the position along the front. Otherwise a front speed consistent with the speed of the low or a default value is used.

3.2.3 High Pressure Centers

Because the names and the location of high pressure centers are used to determine the direction of front movement through the "airmass_toward" property described in Section 2 of this report, high pressure centers must be created on each side of every front. Location of highs are determined from gradients if a hull exists and from the orientation of the fronts if it does not.

3.3 Update Analysis

Because synoptic features such as high and low pressure systems evolve relatively slowly over a period of many hours or a day or two, it is more effective to provide synoptic model maintenance rather than continual model re-creation. This is particularly true since the observation stations that report data in the current hour may not be the same stations that reported data in the past hour. Consider the extreme case where a dozen stations reported in the past hour and a front was well positioned within the hull. If only a single station reports during the current hour, it wouldn't be desirable to create a new model based only on that one station.

If an analysis from the prior hour exists, the first action done upon initiation of the analysis is to move all of the meteorological entities based upon their speed and direction of movement. The analysis process then consists of validating and updating each of the entities.

The process of validating each of the meteorological entities consists of checking that it is consistent with the new data. If the entity is consistent with the new observations, nothing is done to it. If there is evidence that the entity exists but is in a somewhat different location, or has a different speed or direction, then these attributes will be adjusted. An example of how this works is illustrated by the movement of a front. If in the synoptic model a previously created front moved passed a station during the current hour but the data from the station indicated that it still hadn't passed, the position of the front has to be moved back.

The process of updating the analysis must also perform two other functions. It must eliminate any meteorological entity that has moved substantially past the region of interest and is no longer relevant to the weather at any of the stations. It must also create new highs, lows, and/or fronts that are indicated by the data.

4.0 FORECASTS

Itasca provides 1-12 hour forecasts for the variables of pressure, temperature, dew point temperature, wind speed and direction, cloud height and amount, ceiling, weather, and visibility at a station location selected by the user. The quality of the forecasts made by *Itasca* is largely dependent upon the quality of the synoptic model generated by the expert system during the analysis process. The forecasts are based upon weather patterns associated with the classical cyclone model modified by current observations, to the extent they are available. This section describes how the forecast process is implemented.

Because of the feedback nature between some of these variables, forecasts of all the variables are made one hour at a time with the prior hour's forecast used as if it were the most recent observation. In other words, the process described below is repeated twelve times.

For each hour of the forecast cycle, a synoptic map is created with the meteorological entities moved to where they are expected to be at the validation time of the forecast. Each variable that is forecast has expected values and expected value changes depending upon the location of the forecast station relative to the pressure systems and fronts. If only a single station is available, the forecast will depend only on the observed data at that station and the classical model of cyclone weather. If observation data are available from other stations, that information is used to improve on the standard forecast. For example, precipitation may normally be forecast to occur near and behind a cold front. If data are available at upstream stations and they show rain occurred several hours ahead of the cold front, the forecast will be modified to show an increase in probability of precipitation ahead of the cold front at the forecast station. Observation data that are able to characterize an approaching airmass will be used when available to improve the post-frontal forecast for the forecast station. The following paragraphs briefly outline some of the considerations used in forecasting each of the variables.

The first two variables forecast are pressure and wind direction. Forecasts of these variables depend primarily on the movement and location of the pressure systems and fronts. Pressure is forecast by evaluating the change in position of the surrounding pressure systems relative to the forecast station. A preliminary wind direction is forecast by evaluating the location of the surrounding pressure systems and front. This wind direction may then be modified by diurnal effects such as sea-breeze/land-breeze, vertical mixing during the convective portion of the day, or by topography.

Wind speed is the next forecast variable. Wind speed often exhibits a diurnal cycle, particularly near high pressure centers and on clear, dry nights. Daytime wind speeds are affected by vertical convective mixing and by local terrain and local circulations. Wind speeds tend to be at least 12-15 knots and gusty for about 24 hours after a cold frontal passage. Basic wind speeds tend to be similar within the same airmass.

Temperature also normally exhibits a diurnal cycle, but this cycle is damped by cloud cover and a humid atmosphere. Temperature advection ahead of a warm front or behind a cold front can eliminate or reverse the diurnal cycle. Liquid precipitation can result in substantial cooling.

The dew point temperature tends to change due to advection of ahead of a warm front and behind a cold front. Otherwise it will remain relatively constant.

Advective changes in cloud cover often occur in predictable ways with the advancing and receding fronts. High clouds tend to be very persistent. Middle clouds when they occur alone may be quite patchy. Middle clouds will increase and lower with the advancement of a warm front or low pressure center. Convective clouds will develop at the convective condensation level if the convective temperature is reached. The ceiling is determined mathematically from the cloud forecasts.

Convective precipitation can be forecast using sounding data and stability indices with a consideration of additional lifting or destabilizing mechanisms such as cooling aloft. Stable precipitation is associated with warm fronts which can be enhance by large scale orographic lifting. The shear-of-shear vector may be used to find areas of developing instability. Radiation fog occurs when the atmosphere above the surface is dry, the wind is light, and the temperature decreases to near the dew point temperature, particularly over water surfaces or wet soils. Advective fogs should be forecast if warm humid air flows over colder ground, particularly if the dew point of the moist air is higher than the temperature of the cold air. Evaporation fog develops when cold outbreaks flow over warm water, or with sustained precipitation that cools the air.

5.0 USER INTERFACE

The user interface for *Itasca* was developed using Neuron Data's Open Interface software package, and it consists of several windows in which the user can initiate the analysis and forecast, view displays, or enter data. The user should be able to quickly learn the interface and find it easy to use.

The Monitor window is the primary window in *Itasca*. It consists of a menu bar with drop-down menu choices, current time and status information, a small map showing the locations of observation stations and whether or not they had current data, and a list of stations with their names and location. Operational users will interact with *Itasca* primarily through the Monitor window menu. Figure 3 displays the evaluation version of the menu, though not all of the choices are operational.

Session	Edit	Make	View	Developer
New	Stations...	Analysis	Observations	<i>AutoForecast</i>
Open...	Observations	Set Forecast Station	Surface...	<i>Forecasts to File</i>
Save	Surface...	Forecast	Upper Air...	<i>Transcript to File</i>
Save As...	Radiosonde...		Products	<i>AutoRun</i>
Quit...			Analysis...	<i>1 Hour</i>
About Itasca...			Forecast...	<i>6 Hours</i>
				<i>12 Hours</i>
				<i>Specify...</i>
				<i>Make Forecast Now</i>
				<i>Inspect Objects...</i>

Figure 3: The Monitor Menu

The **Session** menu is used to create new sessions, recall previous sessions, save sessions, and shut down *Itasca*. Saving sessions allows one to exit *Itasca*, do other tasks, and then recall the session to enter more data, make analyses, and make forecasts using data retained from the session up to the point it was saved. Only the Quit and About *Itasca* items are active in the evaluation system.

The **Edit** menu items activate station and observation editors. If data are to be entered manually during operational use, the Station Editor would be invoked first. The user would be required to supply information about the stations used during the session. Once the station data are entered, the user would then select the surface or upper-air editor and enter the appropriate observation data for the first hour. For each subsequent hour, the user must enter all of the new observation data.

This **Make** menu is used to produce analyses and forecasts. After the hour's data are entered, the user initiates the analysis by selecting **Analysis** in the drop-down menu. If not already selected, the user next selects the **Set Forecast Station** and specifies which station will be the forecast site. Lastly, the user selects **Forecast** to have the system produce the 1-12 hour forecast.

The **View** menu is used to activate observation and product viewers. The observation and product viewers are self-contained windows with menus and buttons that permit the user to choose how they view the data. Data and observations are available in a graphical format wherever possible to aid in understanding and interpretation. The **Surface** data viewer allows the user to specify one or two stations and view time series of observation values from the station(s) selected. The user may select whether to view these time series in graphical or text form. The **Upper Air** viewer allows users to display soundings in either Stüve or skew-T format. Soundings from up to two different stations or times may be overlaid. A variety of derived quantities from the sounding data are provided, including basic thermodynamic values and stability indices. The user may compute the wind shear or the shear of shears at any arbitrary level in the sounding.

Selecting the **Analysis** viewer under the **View** menu item opens another window that provides a geographical background map of the observation and forecast region. The user can make choices from the menus in this window on what to view. On this analysis map, this viewer can depict any of the following:

- Synoptic elements (pressure systems, fronts)
- Isobars, isotherms and other fields, including time changes
- Station models
- Current and past weather maps
- Surface and upper-air fields

Fields in the viewer are updated automatically every time an analysis is made. The **Forecast** viewer provides a time-series view of the (up to) 12 hours of observations and the latest 1-12 hour forecast. In appearance, this display is similar to the surface observation viewer. The user can select to view the forecast in either graphics or text.

The **Developer** menu is not part of the operational version of *Itasca*. It is included in the evaluation system to allow processing of archived data.

6.0 SUMMARY

This report describes the construction and operation of the knowledge-based mid-latitude forecasting system *Itasca*. *Itasca* is intended to be a system that could provide "first-in" capability or in regions where data may be sparse. To this end, *Itasca* is designed to incorporate an arbitrary and possibly changing number of reporting stations. It is assumed that data may not be available at every hour for every station.

Itasca simulates the methodology of experienced forecasters by first assimilating all data available into a synoptic model of the meteorological environment. A forecast of several meteorological values may then be initiated. The forecast station may be any of the stations available, but data at the forecast station is assumed to be hourly and always available. The user interface of *Itasca* is easy to learn and easy to use. It permits viewing the synoptic maps in graphical form, and it allows the presentation of observation data and forecast data in graphical form.